

This Page Is Inserted by IFW Operations  
and is not a part of the Official Record

## **BEST AVAILABLE IMAGES**

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

**IMAGES ARE BEST AVAILABLE COPY.**

**As rescanning documents *will not* correct images,  
please do not report the images to the  
Image Problem Mailbox.**



## INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification <sup>5</sup> :

H04N 1/32

A2

(11) International Publication Number:

WO 90/05422

(43) International Publication Date :

17 May 1990 (17.05.90)

(21) International Application Number: PCT/US89/05017

(22) International Filing Date: 7 November 1989 (07.11.89)

(30) Priority data:

268,781

8 November 1988 (08.11.88) US

(71) Applicant: MASSACHUSETTS GENERAL HOSPITAL  
[US/US]; 55 Fruit Street, Boston, MA 02114 (US).(72) Inventors: TAAFFE, James, L. ; 74 Appleton Street, Ar-  
lington, MA 02174 (US). KALDIS, Maria ; 9 Hawthorne  
Place, 4M, Boston, MA 02114 (US).(74) Agents: SMITH, James, M. et al.; Hamilton, Brook, Smith  
& Reynolds, Two Militia Drive, Lexington, MA 02173  
(US).(81) Designated States: AT (European patent), BE (European  
patent), CH (European patent), DE (European patent)  
FR (European patent), GB (European patent), IT (Euro-  
pean patent), JP, LU (European patent), NL (European  
patent), SE (European patent).


## Published

*Without international search report and to be republished  
upon receipt of that report.*

(54) Title: DIGITAL IMAGE PROCESSING AND DISPLAY SYSTEM

## (57) Abstract

An image processing and display system employs a high speed image LAN (31) for transferring image data separate from control data. A general purpose LAN (21) carries graphics data and control data of the system. One more image display controller (25) than display monitor (23) is used to provide off-line processing of images prior to display of the images. An analog cross-bar switch (45) provides the necessary connections between image display controllers (25) and the display monitors (23) to provide the desired views of selected images. Images are initially provided to the system from an image archive in a compressed manner. The images are decompressed and held in a local cache (35). With respect to each word of image data, different bytes of the word are compressed by different compression schemes to optimize transfer time of image data. A method is employed for transparently sharing a high performance image archive processor among various workstations of display monitors (23).



*FOR THE PURPOSES OF INFORMATION ONLY*

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	ES	Spain	MG	Madagascar
AU	Australia	FI	Finland	ML	Mali
BB	Barbados	FR	France	MR	Mauritania
BE	Belgium	GA	Gabon	MW	Malawi
BF	Burkina Faso	GB	United Kingdom	NL	Netherlands
BG	Bulgaria	HU	Hungary	NO	Norway
BJ	Benin	IT	Italy	RO	Romania
BR	Brazil	JP	Japan	SD	Sudan
CA	Canada	KP	Democratic People's Republic of Korea	SE	Sweden
CF	Central African Republic	KR	Republic of Korea	SN	Senegal
CG	Congo	LI	Liechtenstein	SU	Soviet Union
CH	Switzerland	LK	Sri Lanka	TD	Chad
CM	Cameroon	LU	Luxembourg	TG	Togo
DE	Germany, Federal Republic of	MC	Monaco	US	United States of America
DK	Denmark				

-1-

DIGITAL IMAGE PROCESSING AND DISPLAY SYSTEMBackground of the Invention

A typical computer network system includes a digital processor, a main disk for storage and several workstations which are serviced by the digital processor and which share the information stored in the main disk. Each workstation generally includes one or more display and keyboard terminals which enable a user to process and view data retrieved on a selected basis from the main disk. In such workstations the cost of the display hardware and the speed with which images can be loaded from the main storage or an image archive is critical and of prime importance. Nowhere is this more so than in a workstation with many high resolution displays, such as one would find in a medical workstation or other workstation used for diagnostic purposes and in an image processing workstation.

In the transmission of image data from an image archive or an image processor to a display workstation, the speed at which the image is transmitted is of paramount importance. There are generally two solutions that can be addressed to decrease this time. One solution is to increase the bandwidth of the communication channel between the image archive and workstation. A second solution is to compress the data prior to transmission. Data compression, in this context, is only effective if the decompression time is negligible in relation to

-2-

the time saved to transmit the data. There are two main classes of image compression algorithms, noisy algorithms and exact algorithms. Noisy algorithms are those that produce a small difference between the original image and one that has undergone a compression-decompression cycle. Exact algorithms, are those that leave the image completely unchanged upon being decompressed.

Various arrangements of display hardware and use of various compression schemes have been used to provide optimum image display systems within a computer network assembly but none are problem free, simple or inexpensive.

#### Summary of the Invention

The present invention provides a digital image processing and display system which optimizes the time and cost efficiencies of display hardware and data transmission. One feature of the system is that separate communication channels are provided to workstations of the display system from a digital processor and image archive for control data and image data. In the preferred embodiment, the image archive and digital processor are tightly coupled forming an image archive/processor unit. A one way optical LAN (Local Area Network) carries image data from the image archive/processor unit to a workstation of the display system while a general purpose LAN carries control data, image address data and user supplied graphics data from the unit to a

30

-3-

workstation. Within the workstation, a high-speed-image bus is used to carry image data alone while a general-purpose-control bus is employed to carry control and graphics data.

5 In addition, the digital image processing and display system includes a display controller scheme in which a group of display monitors is driven in a paint-free fashion, that is by an abrupt, whole image changeover as opposed to a line by line  
10 changeover of images, such that instantaneous changeover from one image to a new image is provided upon the display monitors. In the display controller scheme, there is at least one more controller than monitor, and each controller is  
15 connected to a different monitor to drive the monitor. Each controller is also able to be loaded with a different image to display that image on the respective monitor.

In the preferred embodiment, a cross-bar switch  
20 is used to connect the controllers with the monitors in a one-to-one correspondence relationship, there being one controller at all times unconnected to any monitor. The unconnected display controller is free to be supplied with image data from the image  
25 archive through a local buffer or from another display controller which is connected to a monitor. Once the unconnected display controller is loaded with the desired image data, a currently connected display controller is disconnected from its monitor  
30 while the unconnected display controller is

-4-

connected to the monitor. The unconnecting of the one display controller and connecting of the other display controller holding the desired new image data provides an instantaneous changeover between  
5 images, and as a result, a paint-free, steady display of a succession of images on the monitor.

With respect to the aspect of minimizing data transfer time between the image archive/processor unit and the workstations, the present invention  
10 adaptively applies different ones of various compression schemes to each part of a data word of an image according to high order and low order bits of the word. That is, different compression schemes are used for the high and low order bits of the same  
15 image data word, and the compression schemes of either (or both) the high order or low order bits may change from one sequence of image data words to the next. Such an adaptive application of compression schemes enables those compression  
20 schemes which best condense long series of similar bytes and those compression schemes which more advantageously condense short series of similar bytes to do so in the same image. Typically, prior art systems employ only one compression scheme to  
25 handle all data words of an image. In the preferred embodiment, a Run Length Coding method is used to compress high order bytes of image words, and various Run Length Coding schemes and two different Delta Coding schemes are used to compress the low

30

-5-

order bytes of the image words which tend to change more frequently than the high order bytes.

The foregoing aspects of the display system of the present invention enable "virtual image processing". In such image processing, a network of workstations share the image archive/processor unit without compromising speed of performance. This is accomplished by a processing scheme which does not require image data to be passed between the workstations and image archive/processor unit but only from the image archive/processor unit to the workstations. As a result, the use of the image archive/processor unit by each workstation is transparent to the other users. In a preferred embodiment, any one monitor may temporarily serve as a graphic or iconic control screen showing the placement of images on all other monitors. Furthermore, user interaction with this control screen can effect the movement and selection of both the images being displayed and those to be displayed on the other monitors. The one chosen monitor thus serves as a central control monitor through which the user requests and arranges images to be displayed with respect to the other monitors. For each monitor of the workstation, the user may specify which images and the order in which those images are to be displayed, the order including serial arrangement as well as simultaneous display in different partitions, i.e. rectangular areas, on the monitor. Such a display of the specified images



-6-

is provided without noticeable delay caused by another user at another workstation of display monitors requesting the same or other images to be displayed.

5 In a particular medical application of the image processing and display system, radiological images are selected and viewed from the monitors of a workstation. The images may be displayed in various orientations and gray scale transformations  
10 on several monitors and/or in different partitions of one monitor as desired for diagnosis or other purposes.

#### Brief Description of the Drawings

The foregoing and other objects, features and  
15 advantages of the invention will be apparent from the following more particular description of preferred embodiments of the invention, as illustrated in the accompanying drawings in which like reference characters refer to the same parts  
20 throughout the different views. The drawings are not necessarily to scale, emphasis instead being placed upon illustrating the principles of the invention.

Figure 1a is a schematic view of a network of  
25 workstations connected to a central image archive/processor embodying a display system of the present invention.

-7-

Figure 1b is a schematic view of the image handling and display portion of one workstation of the network of Figure 1a.

Figure 1c is an illustration of an image display format supported by the multiple monitors in the workstation of Figure 1b.

Figure 1d is a schematic view of one portion of the display format of Figure 1c supported by a control monitor.

Figure 2a is a schematic view of a crossbar switch employed in the workstation of Figure 1b.

Figure 2b is an electrical schematic diagram of a transfer gate employed in the crossbar switch of Figure 2a.

Figure 3 is a schematic view of the data structure employed by the Run Length Coding - Long method of compression of image data in the display system of Figures 1a and 1b.

Figure 4 is a schematic view of the data structure employed by the Run Length Coding - Short method of compression in the display system of Figures 1a and 1b.

Figure 5 is a schematic view of the data structure employed by the Delta-2 method of compression in the display system of Figures 1a and 1b.

Figure 6 is a schematic view of the data structure employed by the Delta-4 method of compression in the display system of Figures 1a and 1b.

-8-

Figure 7 is a schematic view of the data structure employed by the Nibble 8/16 method of compression in the display system of Figures 1a and 1b.

5        Figure 8 is a schematic view of the data structure employed by the Nibble 12/16 method of compression in the display system of Figures 1a and 1b.

10       Figure 9 is a schematic view of the data structure employed by the Nested method of compression in the display system of Figures 1a and 1b.

15       Figure 10 is a schematic view of the data structure used to encode one image with the methods of compression of Figures 3-9 adaptively applied to high and low order bytes of data words of the image.

Figure 11 is an illustration of image display software used to support the workstations of the network in Figure 1a.

20       Figures 12a-12d are illustrations of the data structures employed in the software of Figure 11.

#### Detailed Description of the Invention

25       An image/graphics display and processing system embodying the present invention is provided in Figures 1a and 1b. Images of interest are previously generated by other means and are stored in an image archive/processor 17, in a compressed or uncompressed form. The system of the present invention provides retrieval of selected images from

30

-9-

the archive/processor 17 and processing and display of retrieved images, with or without superimposed graphics, at various workstation terminals 15 in a network of workstations which share both a system control processor 13, referred to as the host computer, and the image archive/processor 17. Further, the system provides a view of desired images over several monitors 23 of a workstation 15 and/or over different user partitions of a single monitor in a user specified manner.

In the preferred embodiment, within each workstation 15, one user selected monitor 70 provides a view of the user selected images being displayed and those to be displayed on each remaining monitor 23 of the workstation 15 as illustrated in Figures 1c and 1d. Thus, through this one monitor 70, the user arranges each desired image to be displayed currently or in the future on a specified monitor or partition thereof, at a desired grey level transformation and/or orientation. Graphical and/or iconic indications aid the user in accomplishing his desired ordering of images, such as a miniaturized display or other identification of the next image 72 to be displayed on a monitor 23 and an indication of that relationship between the image 72 and monitor 23. In addition, a single cursor is employed and is moveable from monitor to monitor in a fashion that allows all the monitors to effectively be a single large display screen. Further, the one selected

-10-

monitor 70 serves as an organizational member of the workstation 15 which enables the user to view desired images on the remaining workstation monitors 23 while, in a provided work area 76, processing  
5 another image off line for future viewing. The one monitor 70 may be any of the monitors of the workstation 15 as described later.

With reference to Figure 1a, through a general purpose LAN 21 the host computer 13 controls the  
10 flow of image data and graphics data to the workstations 15. The image data of selected images requested at by a workstation are passed to the workstation 15 on an optical Local Area Network (LAN) 31 which is a high speed, unidirectional,  
15 primarily image carrying channel from the image archive/processor 17 to the workstations 15. Graphics data and text data such as geometric markings (lines, boxes, etc.) provided by the user to be overlaid on image data are transferred on the  
20 general purpose LAN 21. Control data and image address data are also passed to the workstations 15 on the general purpose LAN 21.

For greater time efficiency, the image archive/processor 17 preferably employs a cache 44 (Figure  
25 1a) of the most recently accessed images and an anticipatory cache 46 for holding images determined to be images that are likely to-be-accessed next. These caches 44, 46 operate in a manner common in the art.

30

-11-

In addition, to ensure fair utilization of image archive/processor 17 by many workstations 15, each command asked of image archive/processor 17 has associated with it a priority. This prevents those  
5 image processing commands which take inordinate time and resources from blocking or delaying more routine commands generated from other workstations. To implement such priority, a fast queue and a slow queue are employed by the image archive/processor  
10 17.

In each workstation 15 as shown in Figure 1b, an image only handling module called an image LAN controller 37 receives the image data directed to that workstation 15 from the Optical LAN 31. The  
15 image LAN controller 37 is coupled to the Optical LAN 31 in a receive only manner, and preferably comprises TAXI chips of Advanced Micro Devices, Inc. or the like for receiving the image data. After receiving the image data, the image LAN controller  
20 37 transfers the image data on a 32 bit wide image processing bus 39, or appropriate bus and places the image data into a cache/working buffer 35. In some cases, the image LAN controller 37 decompresses the image data before placing the image data into the  
25 cache/working buffer 35.

The cache/working buffer 35 is a multimegabyte memory whose primary purpose is to locally hold, for the workstation 15, image data representing images retrieved from the image archive/processor 17.  
30 Preferably the cache/working buffer 35 is formed of

-12-

standard dynamic RAM chips with access times of about 100 nanoseconds. The cache/working buffer 35 is dual-ported and is accessed sequentially through either the image processing bus 39 or a suitably wide (e.g. 32 bit wide) high speed image bus 19. More specifically, the image processor bus 39 conveys image data between the cache/working buffer 35 and the image graphic processors, and the high speed image bus 19 transfers image data from cache/working buffer 35 to image display controllers 25 (to be described). Between the cache/working buffer 35 and the high speed image bus 19 are a lookup table 41 and a subsampling subsystem 42. The lookup table 41 performs grey scale transformation by simple substitution through a RAM which is preferably 4K bytes deep. That is, look-up table 41 translates image data from an input scale of pixel values to a desired output scale of pixel values on their way onto the high speed image bus 19. The contents of look up table 41 are modifiable by the image-graphic processor 43.

Subsystem 42 effects the magnification or minification of images to be displayed. The subsystem 42, in the case of magnification, duplicates the image data flowing therethrough and in the case of minification subsamples the image data flowing therethrough. For example, for an image to be magnified by a factor of 2, the subsystem 42 replicates each pixel and replicates each line of the image.

-13-

Also within each workstation terminal 15, there is a display sub-system which enables the workstation 15 to display and manipulate images and graphics on a predetermined number of monitors 23 or  
5 screens connected to the workstation terminal. An illustration of a single workstation 15 with four display monitors 23 is provided in Figure 1b although any number of monitors of any type could be used. Preferably, each monitor 23 displays images  
10 in black and white and graphics in color. In some cases color images and black and white graphics may be preferred and are similarly provided. This is accomplished through an assembly of video mixers 29, graphic display controllers 27 and image display  
15 controllers 25. There is one video mixer 29 and one graphic display controller 27 for each monitor 23, and at least one more image display controller 25 than number of monitors 23. Thus, each monitor 23 is connected to a different video mixer 29 which,  
20 combines video graphic signals from one dedicated graphic display controller 27 and video image signals from one changeable image display controller 25, to be described later.

Each graphic display controller 27 is a 768 by  
25 1024 by 4 bit display controller whose primary purpose is to generate a high bandwidth analog video graphic display signal from digital information. The display memory of each graphic display controller 27 supports limited access, simultaneous  
30 with display, via a general purpose control bus 20.



-14-

Also, each graphic display controller 27 contains a triple 4-bit RAMDAC and thus, outputs a video RGB signal in the case of color graphics. Other output signals of each graphic display controller 27 are  
5 attribute bits for indicating graphic blinking and presence of graphic data. Modification of the graphic display controller 27 memory and creation of graphics through the graphic display controller 27 is performed by the image graphic processor 43  
10 through the general purpose control bus 20.

Each image display controller 25 is a 1536 x 2048 x 8 bit display controller whose primary purpose is to generate a high bandwidth, for example approximately 250mhz, analog video display signal  
15 from a digital pixel map. The display memory on each image display controller 25 does not support simultaneous access but instead can be taken off line (i.e. disconnected from video mixer access) and accessed (i.e. read from or written to) via the high  
20 speed image bus 19. Because image data at this stage has already been translated by look-up table 41, each image display controller 25 does not depend on a look-up table of its own unlike display controllers in existing display systems, and each  
25 display controller 25 outputs a single standard analog video signal plus synchronization signals. Generally, each image display controller 25 comprises 3 megabytes of RAM, a rectangle loader (described later), a display timing generator, a

30

-15-

video digital-to-analog converter and a high speed image bus interface.

The high speed image bus 19 provides an image-data-only transfer channel from the cache/working  
5 buffer 35 to the image display controllers 25 and between image display controllers 25. In the preferred embodiment, the high speed image bus 19 carries four 8-bit image pixel values per transfer and normally operates in a block-transfer burst  
10 mode. On the other hand, the general purpose control bus 20 is used to carry control and status information to all components of the workstation 15 and does not necessarily carry image data.

Having one more image display controller 25  
15 than monitor 23 enables the display subsystem of workstation 15 to provide the user with a more paint-free (i.e. abrupt changeover) looking display from image to image. Further, the easiest and fastest method of loading an image display  
20 controller 25 is to do so when the image display controller 25 is not currently supporting a monitor 23 and is totally available for receiving from the high speed image bus 19 a newly retrieved or modified image to be displayed next. A crosspoint  
25 or crossbar switch 25 is used to quickly connect the disconnected, available image display controller 25 with any monitor 23 and thus route the video image signals from the available image display controller 25 to any monitor 23. This connecting of any four  
30 of the five image display controllers 25 to the

-16-

monitors 23 is schematically illustrated in Figure 1b and shown in more detail in Figures 2a and 2b.

In a preferred embodiment, the crossbar switch 45 is a configuration of transfer gates 47 positioned at each intersection of the five input rows and four output columns, the rows being input lines (Ein) from the image display controllers 25 and the columns being output lines (Eout) to the monitors 23. Only one transfer gate 47 per column is allowed to be conducting at one time. Each transfer gate 47 performs as a high speed analog switch and conveys high frequency signals with low distortion and minimal noise. Referring to Figure 2b within each transfer gate 47, a diode 51 is forward biased by an incoming control signal 49 and serves as a low resistance ohmic device which carries the input signal (Ein) unchanged to the output line (Eout). When the diode 51 is reverse biased by the control signal 49, the high impedance of the diode 51 effectively blocks the input signal (Ein). Capacitors C1 and C2 perform DC blocking functions and have values chosen sufficiently large so as to not degrade the low frequency components of the transferred input signal. Resistor pairs R1-R2 and R3-R4 perform impedance matching functions.

An example of the operation of the crossbar switch 45 during the loading of one image onto a full display screen of monitor 23 labelled Monitor<sub>2</sub> in Figure 1b follows. Assume that crossbar switch currently directs image signals from input line Ein

-17-

1 supplied by image display controller 25 to  
Monitor<sub>2</sub> through output line Eout 2. The image  
graphic processor 43 (Figure 1b) determines which  
image display controller 25 is the available one.  
5 Assuming the available image display controller 25  
is the one which supplies input line Ein 3 (Figure  
2a), the image graphic processor 43 transfers the  
desired image data from the cache/working buffer 35,  
across the high speed image bus 19 and into the  
10 image display controller associated with input line  
Ein 3. After the controller is fully loaded, the  
image graphic processor 43 generates a control  
signal for reverse biasing diode 51 at all transfer  
gates 47 in the second column (i.e. the Eout 2 line  
15 connected to Monitor<sub>2</sub>) of Figure 2a except the one  
located in the row labelled Ein 3. Consequently,  
the crossbar switch 45 disconnects the initial  
connection from Ein 1 to Eout 2 (i.e. from image  
display controller 25 at Ein 1 to Monitor<sub>2</sub> at Eout  
20 2) and routes only the video image signal from the  
image display controller Ein 3 to Monitor<sub>2</sub>.

It is noted that the foregoing load is  
accomplished at the highest rate that the cache  
working buffer 35 can sustain since the available  
25 image display controller's complete memory bandwidth  
is available. Further, from the user's point of  
view, the newly loaded image just popped up on the  
monitor screen as opposed to being painted line by  
line on the screen.

30

-18-

Also, it is understood that color images having red, green and blue components provide 3 inputs per input line shown in Figure 2a. Hence, a crossbar switch for handling RGB image data may be used in  
5 the case of color images.

Referring back to Figure 1b, a general purpose LAN interface 33 and the image graphic processor 43 provide the decision making and driving forces of the image/graphics display and processing system  
10 within a workstation 15. The image graphic processor 43 coordinates all data transfers between the various components within a workstation 15, performs all local image processing, and performs all graphic generation, and coordinates interaction  
15 with the user. In particular, the image graphic processor 43 performs user desired local image processing on requested images held in the cache/working buffer 35 and transfers the processed images to image display controllers 25 which  
20 subsequently drive monitors 23 with the processed images.

Further, the image-graphic processor 43 in tandem with the cache working buffer 35 enables the cache working buffer 35 to subsample an image if the  
25 partition of a monitor screen (the desired destination) is smaller than the image. In that situation, the image-graphic processor 43 also specifies to the receiving image display controller 25 the first column on the monitor screen to start  
30 laying in the image, the number of pixels to copy

-19-

before proceeding to the next row in the bit map for the screen partition and other pertinent data moving information. Each image display controller 25 in turn employs an image rectangle loader 38 for  
5 receiving and interpreting the specified information from cache working buffer 35. The image rectangle loader 38 generates the proper digital pixel image and provides this map to the associated DAC which subsequently drives the target monitor partician  
10 with the desired image.

Preferably the image graphic processor 43 comprises either a processor of the TI34010 type manufactured by Texas Instruments, Inc. including its attendant program and scratch memory or an  
15 MC68030 by Motorola and its attendant program and scratch memory. Other image graphic processors which similarly provide such processing and moving of image data are suitable.

The general purpose LAN interface 33 carries to  
20 the image archive/processor 17, via the general purpose LAN 21, information about the workstation 15, such as, image display and processing commands, and general system coordination commands. Since all command and control information travels over the  
25 general purpose LAN 21, the optical LAN 31 is left available to carry image data alone. That is, the optical LAN 31 is able to transfer image data free of control data and therefore achieve an image transfer efficiency rate of about 95%. A useable  
30 image data transfer rate of about 13 megabytes per

-20-

second is also achieved unlike image transfer rates of existing display systems.

Preferably, the general purpose LAN interface 33 is an Ethernet link.

5       The serial connection of the Optical LAN 31, Image LAN controller 37 and cache/working buffer 35 through a dedicated image processing bus 39 provides for a cost efficient system. The cost of implementing a fully functional Optical LAN  
10 controller (i.e. one which is responsible for both image buffer memory and the handling of data control signals) is about ten times more expensive than using the image LAN controller 37 for buffer memory and the separate general purpose control bus 20 for  
15 control purposes. In addition, the image LAN controller 37 directly providing image data to the cache/working buffer 35 from the Optical LAN 31 through image processing bus 39, requires no permanently designed packet size for the transfer of  
20 data. Thus the packet size can be adjusted dynamically to be most efficient for each type of image data transferred. Furthermore, there are no intermediary data packet handling modules from which the image data must be copied before the image data  
25 is copied into the cache/working buffer 35. Hence the present invention omits unnecessary copying of image data found in other display systems between the Optical LAN 31 and cache/working buffer 35.

Time and economy advantages in image retrieval  
30 and display of images within a workstation 15 are

-21-

also gained by using the separate buses (i.e. the high speed image and general purpose control buses 19, 20) for image data and control data as opposed to one bus within the workstation as is commonly  
5 done. This point is most easily understood by way of the following examples of operation of the image/graphics display and processing system of Figure 1b.

When a user of a workstation 15 requests a  
10 number of images to be loaded into the workstation, the image-graphic processor 43 forwards the request to the image archive/processor 17 over the general purpose LAN 21. In response, the image  
archive/processor 17 provides the image graphic  
15 processor 43 pertinent specifications of the images and an indication of whether the images will be shortly transmitted on the optical LAN 31. The image graphic processor 43 through the general purpose control bus 20 allocates room in the  
20 cache/working buffer 35 for the requested images. Currently held image data within the cache/working buffer 35 is deleted as space is required. The image graphic processor 43 through the general purpose control bus 20 initializes the image LAN  
25 controller 37 to receive and possibly decompress the requested images. The image-graphic processor 43 of each workstation 15 and the image archive/processor 17 communicate with each other to provide the requested images to the proper workstation 15 and to

30



-22-

satisfy any other requests from other workstations  
15.

The images are packetized and sent one at a  
time from the image archive/processor 17 to the  
5 image LAN controller 37 on the optical LAN 31.  
During the transfer of requested images, the  
image-graphic processor 43 monitors the status of  
the image LAN controller 37 and cache/working buffer  
35 to assure that the image transmission was error  
10 free. If not, the image-graphic processor 43  
requests the image archive/processor 17 to  
retransmit the needed image data.

Subsequent to the retrieval of the requested  
images, a user request to display on a particular  
15 monitor 23 one of the retrieved images is satisfied  
in the following manner. The image graphic  
processor 43 causes (i) the desired image to be  
found in the cache/working buffer 35, (ii) the image  
data representing the image to be translated by the  
20 lookup table 41 and possibly modified by subsample  
subsystem 42, and (iii) the translated/modified  
image data to be placed on the high speed image bus  
19. The image data is transferred through the high  
speed image bus 19 to the image display controller  
25 which is currently available as detected by the  
image-graphic processor 43 over the general purpose  
control bus 20. As explained in detail above, one  
of the several image display controllers 25 is  
always available. Through the general purpose  
30 control bus 20, the image graphic processor 43

-23-

enables the available image display controller 25 and the crossbar switch 45 to drive the particular monitor 23 with the desired image data.

It is noted that all control data was carried  
5 on the general purpose control bus 20, and as a result the high speed image bus 19 was enabled to carry image data alone. Further, the highest rate of data transfer from the cache/working buffer 35 to the image display controller 25 is achieved by the  
10 available image display controller 25 having its complete memory bandwidth available because it is not currently driving a display monitor 23.

When a user requests a grey scale transformation of a displayed image on a monitor 23  
15 the following occurs. Over the general purpose bus 20, the image-graphic-processor 43 checks the cache working buffer 35 for the stored copy of the one displayed image. If a copy of the one image is not found, the image graphic processor 43 requests over  
20 the general purpose LAN 21 the image from the image archive/processor 17. Once the copy of the image is found in the cache/working buffer 35, the image graphic processor 43 locates over the general purpose bus 20 the available image display  
25 controller 25. Next, through the general purpose control bus 20, the image graphic processor 43 generates a new lookup table 41 in the cache working buffer 35 according to the desired requested grey scale transformation. The image graphic processor  
30 43 transfers the copy of the image stored in the

-24-

cache/working buffer 35 from the buffer, through the new lookup table 41 and optionally subsystem 42, across the high speed bus 19 and into the available image display controller 25. Through the crossbar switch 45, the image graphic processor 43 enables the newly loaded available image display controller 25 to drive the same monitor 23 on which the image was formerly being displayed but with a grey scale transformation as requested.

10        If the image which is desired to be processed is currently displayed on only one partition of a monitor 23 (as opposed to a full screen display) and other images which are not to be changed are displayed in the remaining partition of the monitor 23, then the other images must be preserved in the image display controller switching process described above. This is accomplished in the following manner. After the image graphic processor 43 has located the available image display controller 25, 15 the image graphic processor copies the full contents of the image display controller 25, which is currently enabling display of the one desired-for-processing image and the other images, into the available image display controller 25. This copy is 20 accomplished at full memory bandwidth across the high speed image bus 19 and is performed as the data is addressed for display without taking the currently used image display controller 25 off line. The copy time is at most the amount of time to 25 display two or three image frames. This type of 30

-25-

data transfer over the high speed image bus 19 from one image display controller 25 to another is called "snapping". The image graphic processor 43 then proceeds to generate the new look-up table 41 and subsequently transfers the copy of the to-be-processed image from buffer 35, through new look-up table 41, optionally through subsampler system 42, across high speed bus 19 and into the proper partition (i.e. the position of the subject image) of the available image display controller 25. Switching of crossbar switch 45 to drive the monitor 23 with the newly loaded image display controller 25 is then as previously described.

In a request to edge enhance an image, the image graphic processor 43 performs the calculations required and properly modifies the desired image. Through the general purpose control bus 20, the image graphic processor 43 allocates sufficient space in the cache working buffer 35 to temporarily hold the edge enhanced image. Throughout the edge enhancement process, the image graphic processor 43 uses the image processing bus 39 to transfer image data processed by the processor 43 to and from the cache/working buffer 35. In order to display the edge enhanced image, the image graphic processor 43 locates the available image display controller 25 over the general purpose control bus 20 and copies the edge enhanced image to the available image display controller 25 over the high speed image bus 19. The newly loaded available image display

-26-

controller 25 is then connected to a user selected monitor 23 via crossbar switch 45 to drive that monitor with the edge enhanced image.

5 If the edge enhanced image is desired to be displayed on a monitor 23 along with images which are currently being displayed on that monitor 23, then the contents of the image display controller 25 currently connected to the monitor 23 must be snapped into the available image display controller  
10 25 before the edge enhanced image is copied into the desired portion of the image display controller 25 as described above.

A user request to make a duplicate of a displayed image from one monitor 23 or partition  
15 thereof to another monitor 23 or partition thereof is handled by the image graphic processor 43 as a general load. A copy of the desired image is obtained from the cache/working buffer 35 and transferred on the high speed image bus 19 to the  
20 available image display controller 25. Snapping the contents of the image display controller 25 supporting the other monitor 23 into the available image display controller 25 is done as necessary. The newly loaded available image display controller  
25 25 is then connected to the other monitor 23 to display a duplicate of the desired image in the desired manner (i.e. on a full screen or partition thereof).

To further minimize transfer time of image  
30 data, the image data is passed from the image

-27-

archive 17 to a workstation 15 in a compressed form. Generally, images are compressed subsequent to entry into the image archive/processor 17 for storage.

5 Data compression in this context is only effective if the decompression time is negligible in relation to the time saved to transmit the data.

Furthermore, in many applications it is desirable to be able to decompress the image data in real time, specifically, to decompress the image data  
10 simultaneously with the receipt of the data. This implies that the decompression time per unit of data be less than or equal to the transfer time of the unit of data divided by the compression ratio (defined later). For high speed optical LANs (Local  
15 Area Networks) whose unit of data transfer rates approach the maximum hardware speed used to implement them, a further implication directs that the decompression scheme require only a few primitive mathematical operations (i.e. addition,  
20 shifting, etc.) per image pixel.

The foregoing requirements of a time efficient decompression scheme forbid use of the traditional schemes such as Huffman coding, Linear-predictor coding, etc. since they require many complex  
25 operations (multiplication and variable code word size lookup) to be performed for each decompressed pixel. Two classes of compression schemes that satisfy the foregoing requirements are Run Length Coding and Delta Coding schemes. Improved versions  
30 of these compression schemes plus a "Nested Adaptive

-28-

Coding" compression scheme are employed in the present invention.

Generally, each pixel of an image is represented by either:

- 5            sixteen data bits, stored in four  
consecutive nibbles (4 bits) of a sixteen bit  
word;
- twelve data bits, stored in three low  
order nibbles of a sixteen bit word with the  
10          high order nibble set equal to zero;
- eight data bits, stored in two low order  
nibbles of a sixteen bit word with the two high  
order nibbles set equal to zero; or
- 15          eight data bits, stored in an eight-bit  
byte.

Each of the foregoing representations of an image pixel is generally referred to as a "data word".

The "data words" are stored in a file and are usually ordered in a row-by-row fashion such that  
20          the data word representing the upper left hand pixel  
of an image is at the beginning of the file, the  
next data word of the file represents the pixel to  
the right of the upper left hand pixel in the same  
row, and so on with the data word representing the  
25          pixel in the lower right hand corner of the image  
being last in the file. The compression  
manipulations of the present invention do not  
require a particular ordering of the data words in  
an image file corresponding to pixels of an image,

30

-29-

either row-column or column-row ordering will yield essentially the same results.

In the preferred embodiment, four methods of compression are made available for compressing 8-bit (1 byte) long image data and three other compression methods are made available for compressing image data which are longer than 8 bits. The first two compression methods are based on Run Length Coding methods which are applicable for image data in which the more significant bits of data words corresponding to a succession of pixels of an image change less often than the less significant bits of the data words, from one pixel to the next. Such image data is considered to have a high auto-correlation coefficient. The Run Length Coding methods are most advantageously applied to the high order bits of the data words of such image data. The second two compression methods are based upon Delta Coding techniques which are applicable to image data with a medium to low auto-correlation coefficient and local pixel value variance. The next two methods are based on simple nibble packing and are applicable for desired instances of quick minimal compression. The last method involves a technique of nested coding and is a fusion of the techniques for 8-bit data compression with Run Length Coding for the high order bits of image data.



-30-

Compression Method 1: Run Length Coding-Long

In encoding image data with Run Length Coding methods, data is determined to have one of two forms, a run of sames and a run of differents. In a  
5 run of sames, there is a series of similar valued pixels. The corresponding coding comprises the pixel value which is repeated and a count of the number of times the pixel value is repeated. In a  
10 run of differents, there is a series of different valued pixels. The corresponding coding comprises the count of different bytes and the list of the pixel values. Compression Method 1 is efficient for image data having long series or runs of similar pixel values (i.e. where the length of the run is  
15 greater than about 8). The method is most easily understood by the data structure it generates. Such a structure is referred to as a compression unit and is illustrated in Figure 3. Other compression methods generate other compression units which are  
20 discussed later.

The first byte in compression unit 53 is a control byte 55 with three fields. The first field  
59 comprises the two highest order bits of the control byte 55 and provides an indication of the  
25 type of compression method used to compress the image data within the compression unit. The two bits set to zero indicates the Run Length Coding-Long method. As will be described later, different compression methods are used within the  
30 same image and thus at decompression time the type of compression method used for the data within the

-31-

compression unit can be readily determined from the first field of the control byte. The second field is the third highest order bit 57 and provides an indication of the form (i.e. type of run) of the image data being coded. The form bit 57 is set to zero if a run of sames is being coded and is set to one if a run of differents is being coded. The third field, called the count field 61, comprises the remaining five bits of the control byte 55 and indicates the number of bytes of image data in the run indicated by the form bit 57. In the preferred embodiment, for a run of sames the value of the count field 61 is in the range of 2 to 33; for a run of differents the value of the count field 61 is in the range of 1 to 32. Also, biased unsigned integers are used for the values of the count field 61.

The following bytes 65 in the compression unit 53 are the bytes of image data. The values of these data bytes 65 are set to the pixel values which upon decompression are to be repeated in the manner indicated by the form bit 57. For a run of sames, the first data byte 63 will be the pixel value to be repeated the number of times indicated by the count field. For a run of differents, data bytes 1 through data byte n, where n equals the number indicated by the count field 61, hold the pixel values to be used upon decompression.

In order to determine how well a compression method compresses data, the compression ratio  $c$  is

-32-

calculated and equals the ratio of the number of bytes of image data before compression to the number of bytes in the compression unit. For the Run Length Coding-Long method the compression ratio  $c$  for a run of sames equals  $n/2$ , where  $n$  is the length of the run (in bytes). For a run of differents, the compression ratio is  $c = n/(n + 1)$ , always less than unity.

Examples of the compression of data and the resulting compression ratio using the Run Length Coding-Long method follow. An example of Run Length Coding-Long (Sames) is provided first and an example of Run Length Coding-Long (Diffs) is provided second.

Example 1. A run of 16 bytes, each having a value equal to 6.

The compression unit would have:

Control byte = [00, 0, 16]

Data byte 1 = [6]

The compression ratio is  $16/2 = 8$ .

If the length of the run is greater than the maximum value that the count field 61 can represent, the data sequence would be encoded with a series of "sames" compression units.

Example 2. A sequence of bytes having the values 12, 33, 34, 45, 46, 45.

The compression unit would have:

Control byte = [00, 1, 6]

Data byte 1 = [12]

Data byte 2 = [33]

-33-

Data byte 3 = [34]

Data byte 4 = [45]

Data byte 5 = [46]

Data byte 6 = [45]

The compression ratio is  $c = 6/(6 + 1) =$   
5 0.86.

#### Compression Method 2: Run Length Coding - Short

For image data comprised of many short runs (less than 8 bytes per run), the following method is used in which two short runs are coded in one  
10 compression unit. The advantage of this method is that only half a byte is expended for control information per byte of image data whereas in previously described Method 1 a full byte is expended for control information for each data run.

15 The compression unit generated by Method 2 is illustrated in Figure 4. The first byte, the control byte 69, of compression unit 67 has a first field 71 for indicating the compression method used and has two count fields 73, 75 for respectively  
20 indicating the number of bytes of two runs of same encoded in the compression unit. The first field 71 comprises the two highest order bits of the control byte as in Method 1 but are set to 01 to indicate the Run Length Coding - Short method. The two count  
25 fields 73, 75 are each 3 bits long and hold values in the range of 1 to 8.

The byte 77 succeeding the control byte 69 is the data byte which corresponds to the first count

-34-

field 73 and holds the pixel value to be repeated upon decompression the number of times indicated by the first count field 73. Similarly, the data byte 79 succeeding the first data byte 77 holds the pixel value to be repeated upon decompression the number of times indicated by the second count field 75.

For a data sequence which forms a double run of same the compression ratio  $c$  equals  $n/3$ , where  $n$  is the number of bytes of data before compression.

10 An example of a compression unit generated by compression Method 2 for a sequence of nine bytes having the values 3, 3, 3, 3, 4, 4, 4, 4, 4 is as follows:

Control byte = [01, 4, 5]  
15 Data byte 1 = [3]  
Date byte 2 = [4]  
Compression ratio  $c = 9/3 = 3$

#### Compression Method 3: Delta - 2

Generally, in a Delta compression scheme a sequence of pixel values is represented as a sequence of "delta values" which are computed by subtracting a previous pixel value from the succeeding pixel value for each pixel value in the sequence beginning with the second pixel value. For example, the original sequence of 112, 113, 110, 111 is delta compressed to the "difference sequence" of 1, -3, 1. More specifically, the compression unit generated by a typical Delta compression scheme comprises a starting pixel value, the number of

30

-35-

pixel values to be generated from the compression unit upon decompression, and the sequence of delta values.

Further, just as in Run Length Coding compression methods there is a structurally imposed upper limit to the length of a run (the maximum value that the count field could represent), the same behavior manifests here. The number of bits allocated for coding each delta value implies a range of allowable delta values. For instance, 2-bit delta values allow differences from -1 to +2 to be expressed and 4 bits allow ranges from -7 to +8 to be expressed. Thus, long runs of image data which are expressible within a range of delta values may have to be encoded in multiple compression units.

The Delta-2 compression method generates a compression unit 89 illustrated in Figure 5. The first byte, the control byte 81, has a two bit first field 83, which is similar to the first control byte field of the previously described compression methods, and a 6-bit count field 85. The first field 83 is set to 10 to indicate the Delta-2 compression method. The count field 85 is set to the number of delta values held in this compression unit 89. The byte which succeeds the control byte 81 holds the starting pixel value of the uncoded image data (i.e. sequence of pixel values). The remaining bytes 91 each hold four delta values in the range of -1, 0, 1, 2. For a sequence of pixel

-36-

values expressible as a sequence of 2-bit delta values, the compression ratio  $c$  is

$$c = n / (2 + \text{Int}[(n-1)/4] + .5)$$

where  $\text{Int}[ \ ]$  returns the greatest integer.

5       An example of the compression unit generated by compression Method 3 for the sequence of ten bytes having values 3, 3, 3, 3, 4, 3, 5, 6, 8, 9 is as follows:

Control Byte = [10, 9]  
10       starting value byte = [3]  
Delta data 1 = [0, 0, 0, 1]  
Delta data 2 = [-1, 2, 1, 2]  
Delta data 3 = [1, x, x, x]

where  $x$  is an ignored value. The compression ratio  
15        $c = 10 / (2 + 3) = 2$ .

#### Compression Method 4: Delta-4

For some classes of image data, a range of delta values from -1 to +2 is insufficient to efficiently encode the data. Thus it may be  
20       advantageous to use four bits to encode the delta values in a range of -7 to +8. The advantage of such a compression scheme is that for noise data which is inappropriate for Run Length Coding or Delta-2 encoding, a compression ratio of about 2 can  
25       be achieved.

A compression unit 93 for the Delta-4 compression method is illustrated in Figure 6.

-37-

Control byte 95, the first byte, comprises a first field 97 and a count field 99. The two highest order bits of control byte 95 form the first field 97 and are set to 11 to indicate the Delta-4 compression method. The remaining six bits of control byte 95 form the count field 99 and are set to the number of pixel values which are compressed into the compression unit 93.

The second byte 101 of the compression unit 93 holds the starting pixel value. The succeeding bytes 100 each hold up to two delta values in the range of -7 to +8, each delta value being represented in 4 bits. For a series of pixel values expressible as a sequence of 4-bit delta values, the compression ratio  $c$  is

$$c = n / (2 + \text{Int}[(n-1)/2] + .5)$$

where  $\text{Int} [ ]$  returns the greatest integer.

#### Compression Method 5: Nibble 8/16

In the case where 8-bit pixel values are stored in a 16-bit word, the Nibble 8/16 compression method encodes only the meaningful part, that is the lower byte, of the word and discards the upper byte which is always zero. Although this compression method is simple and quick, it is not efficient. The compression ratio  $c$  equals  $2n/(1+n)$ , where  $n$  is the number of pixel values being compressed.



-38-

Figure 7 provides an illustration of the compression unit 14 associated with the Nibble 8/16 compression method. The first byte 16 (i.e. control byte) of the compression unit 14 holds the number of pixel values compressed into the compression unit. The succeeding bytes 18 represent the low order bytes of the subject pixel values.

An example of a compression unit generated by the Nibble 8/16 method compressing the sequence of hex values of 0012, 0023, 0034, 0045, 0045, 0065 follows.

Control byte = [6]  
Data byte 1 = [12]  
Data byte 2 = [23]  
Data byte 3 = [34]  
Data byte 4 = [45]  
Data byte 5 = [45]  
Data byte 6 = [65]

The compression ratio  $c = 2 \times 6 / (1 + 6) = 12 / 7 = 1.71$

#### 20 Compression Method 6: Nibble 12/16

For the case of 12-bit pixel data stored in a 16-bit word, again a simple and quick compression method is to code only the three low order nibbles and not the upper nibble which is always zero.

25 However, compressing in that manner implies that a pixel value will cross byte boundaries (i.e. will be represented by parts of two bytes). Also, although fast, the compression method is not efficient. The compression ratio  $C$  is equal to  $2n / (1 + 1.5n)$ , where  $n$

-39-

is the number of pixel values being compressed. The maximum compression is 1.33.

Provided in Figure 8 is an illustration of the compression unit generated from the Nibble 12/16  
5 compression of hex pixel values 0312, 0345, 0465. The control byte 22, the first byte, of compression unit 78 is set to the value 3, the number of pixel values compressed into the compression unit. The second and succeeding bytes 24 are data bytes each  
10 of which at most hold representations of two nibbles. Preferably, the low order portion of a data byte 24 is set to represent a nibble of a pixel value before the high order portion of the data byte 24 is set to represent the next meaningful nibble  
15 (i.e. non-zero nibble) in the sequence of pixel values. Hence, the lower order bits of the second byte in compression unit 78 is set to a value of 2, the value of the lowest order nibble in the first pixel value 0312, and the higher order bits of the  
20 second byte have a value of 1, the value of the second lowest order nibble in first pixel value 0312. The third byte in compression unit 78 holds in its lower order portion the value 3 which corresponds to the remaining non-zero nibble in  
25 first pixel value 0312. The higher order portion of the third byte holds a 5, the value of the next meaningful nibble in the given sequence of pixel values in order of lowest nibble to highest nibble per pixel value. The rest of the data bytes in  
30 compression unit 78 are likewise set.

-40-

The compression ratio  $c = 2 \times 3 / (1 + 4.5) = 6 / 5.5 = 1.09$ .

#### Compression of Nibble 16/16

For the case of 16-bit pixel values stored in 16 bit words, compression is not efficient. In such a case, the pixel values are held in a compression unit with a control byte and a series of data bytes. The control byte is set to the number of 16-bit pixel values held by the compression unit for a sequence of pixel values. Each data byte holds a byte of a pixel value. Preferably, the first data byte holds the lower order byte of the first pixel value in a sequence of 16-bit pixel values, and the second data byte holds the high order byte of that first pixel value. The third data byte holds the lower order byte of the second pixel value, and so on.

#### Compression Method 7: Nested

As mentioned before in reference to the auto-correlation coefficient, the most significant bits of pixel values in an image will nearly always change very slowly relative to the change of the least significant bits of the pixel values. Hence, a Nested compression method advantageously uses the method of Run Length Coding to compress the high order byte of a 16 bit pixel value or the high order nibble of a 12 bit pixel value, and uses one of the first four described methods (Run Length Coding -

-41-

Long, Run Length Coding - Short, Delta-2, Delta-4)  
to compress the low order byte of the 16-bit or  
12-bit pixel value. In particular, the Nested  
compression method is accomplished by the following  
5 steps:

- (1) Scan the upcoming pixel values to  
determine the run length of constant (i.e.  
same) upper order bytes/nibbles
- (2) If the length of the run is greater than 3  
10 then encode the constant upper order  
byte/nibbles, and apply to the low order  
bytes one of the four mentioned  
compression methods for 8-bit pixel  
values, else use the appropriate nibble  
15 packing compression methods (i.e. Methods  
5 and 6).

An illustration of a compression unit 26 for  
implementing the Nested compression method is  
provided in Figure 9. The first three bytes form a  
20 compression unit header 28. A "type field" of the  
header indicates the major type of compression being  
used. Preferably if the type field is set to 00,  
then the Nested code is indicated and the constant  
field (1 byte long) is set to the value of the high  
25 order byte/nibble which, upon decompression, is to  
be repeated the number of times indicated by the  
field labelled #CU's. The #CU's field also  
indicates the number of following compression units  
30 that encode the corresponding low order bytes  
which have been compressed by one of the compression

-42-

methods for 8-bit pixel values. If the type field of the header 28 is set to 01, then Nibble 8/16, 12/16 or 16/16 Coding is indicated, and the succeeding compression units 30 hold pixel data compressed by one of the foregoing nibble packing methods.

#### Adaptive Compression Methods

Applicant has found that nearly always the auto-correlation coefficient and the variance of the rolling average of pixel values changes for different areas of an image. Thus, to obtain higher compression efficiency the present invention uses different methods on different parts of the image. However, before the time of compression, it is unknown which methods will work best (i.e. will result in the highest compression ratio) in which regions of an image. Hence it is necessary to use adaptive techniques. One adaptive technique is to try each compression method on an upcoming sequence of image data and compare their compression ratios. However, the present invention employs an adaptive technique which checks compression methods in an order according to the length of a run of same in the sequence of image data.

For 8-bit image data, the following table summarizes the order of the methods to try and the number of consecutive bytes with the same value at which no further checking is needed to determine that the method provides optimal compression

30

-43-

relative to the other methods. In addition, the maximum compression ratio is listed.

	<u>Method</u>	<u># bytes to test</u>	<u>Max. Compression</u>
	RLC-Long (Sames)	13	16.50
5	RLC-Short	11	6.33
	Delta-2	8	3.47
	Delta-4	6	1.89
	RLC-Long (Diffs)	NA	0.97

According to the table, if a run of sames of 13  
 10 bytes or more is found in the sequence of pixel  
 values of interest, then the Run Length Coding -  
 Long Method is the optimal compression method  
 because no other method can yield a higher  
 compression ratio. If the run of sames is 11 or 12  
 15 bytes long, however, then the Run Length Coding -  
 Short method provides the optimal compression, and  
 so on. In the case where no run of sames is found,  
 then the pixel values are most efficiently stored  
 sequentially by the Run Length Coding-Long (Diffs)  
 20 Method. No other methods need to be checked, only  
 the length of the run of sames and the table need to  
 be checked. If decisive values can not be reached  
 then the individual compression ratio for the  
 methods are compared and the method having the best  
 25 compression ratio chosen.

For 12-bit and 16-bit image data, a high order  
 byte/nibble run of sames of sufficient length to

-44-

justify Nesting coding must be determined. If such a run of same is not found in the sequence of image data, then the 12-bit or 16-bit pixel values are compressed by one of the nibble packing compression methods described previously. If a run of same which justifies compression by Nesting coding is found, then a determination of the optimal compression method for the lower order bytes is made. Preferably this determination is made by the table described above.

The data structure used to accomplish the foregoing adaptive compression is illustrated in Figure 10. As a result of the determination of an optimal compression method for each series of image data in an image, a sequence of compression units 36 is formed. Each compression unit is one of the types described above, and series of compression units correspond to certain sequences of image data. Each series of compression units 36 is preceded by a compression unit header 34 which indicates the number of compression units 36 before the next compression unit header 34. Other information concerning the associated series of compression units 36 may also be indicated in the compression unit header 34. An image header 32 is used as the first byte of the data structure and indicates that the following sequence of compression unit headers 34 and associated compression units 36 encodes all image data of a particular image. One field in the image header provides the total number of

compression unit headers to be processed for the image. Another field in the image header provides an identification of the image which the data structure represents.

5 In sum, in the present invention, different compression methods are adaptively used for different bytes of image data. In particular, higher ordered bits/bytes are condensed by one of several compression algorithms and lower ordered  
10 bits/bytes are compressed by a different compression method depending on the information carried by the image data word. Other image data words may use even two more different compression methods for their higher and lower ordered bits/bytes  
15 respectively.

Some modes of image data transmission from the image archive/processor 17 to the image LAN controller 37 or a workstation 15 may require the data to be sent in fixed length packets. For  
20 example, to achieve maximum efficiency under the Ethernet LAN (Local Area Network), the packets should be 1,514 bytes long. When these packets are received, they are read into memory buffers of the image LAN controller 37. In the preferred  
25 embodiment, neither a compression unit header or a compression unit span packets and thus memory buffers of controller 37 to avoid spending more time than that which is available to switch input buffers. In the same light, the memory buffers of  
30 image LAN controller 37 may be paged or bank



-46-

selected or of any other type of memory arrangement where more than one physical bank of memory occupies the same logical address and only one bank at a time is enabled for access. As the image data is being  
5 received it is decompressed and placed into the image LAN controller memory buffers. It is preferred that the image data generated upon decompression of each compression unit not span  
10 pages or banks because more time than that which is available to switch banks/pages may be required. Packet boundaries and memory buffer page/bank boundaries may be detected by conventional means and methods to achieve the foregoing.

#### Software Support

15 The foregoing features of the display system of the present invention are implemented by the following image display software routines which form a main support program 80 (Fig. 11) executed in the image graphic processor 43. The image display  
20 software routines are designed to provide (i) "virtual image processing" and (ii) a program which is easy for applications programmers to use. In "virtually" processing images, the image display software processes images locally within a subject  
25 workstation 15 when possible and otherwise processes images through image archive/processor 17 with the assumption that all necessary image data is in image archive 17 (i.e. without transferring image data from workstation 15 to the image archive/processor

17). Such a processing scheme enables each user at a workstation 15 in the network of Figure 1a to retrieve and process images from archive/processor 17 without affecting retrieval and processing performance at the other workstations 15 in the network, and thus in a transparent manner with respect to users of the other workstations 15.

For example, when a user requests an image transform, the support program 80 determines whether the transform is transformable through local processing or is so computationally intensive that the transform must be performed by the image archive/processor 17. In the latter case, the image graphic/processor 43 sends a request over the general purpose LAN 21 to the image archive/processor 17. The request identifies the function to be performed and the image on which the function is to be performed. Additionally, the request indicates the workstation to which the transformed image is to be sent. Upon receiving the request, the image archive/processor 17 performs the indicated function and sends an indication to the specified workstation to request the image when the workstation is ready. The image graphic/processor 43 of the workstation in turn requests the image archive/processor 17 to transmit the transformed image over the optical LAN 31 to the workstation 15. This allows the high performance image processing equipment of the image archive/processor 17 to be transparently shared between the many workstations

-48-

15. For this scheme to work effectively, the image archive/processor 17 preferably contains the image cache 44 which holds for quick access a number of the most recently accessed images.

5       To provide a program which is easy for applications programmers to use, the main support program 80 is designed such that it has few procedure calls for supporting the image handling and displaying of images in each workstation 15. To  
10 that end, the image display software is operated by changing data structures which represent the desired effect instead of by issuing explicit commands to the main support program. Specifically, the image display software manipulates data structures which  
15 represent how the images are to be displayed on the monitors 23 and then invokes a procedure "DO-TRAVERSE-SCREEN-PARTITION-DATABASE" 54 that causes all necessary local and image archive/processing steps to be carrier out.

20       The major data bases of the main support program 80 with their associated data structures as illustrated in Figures 12a-12d are a screen-layout-database 50, an image database 60, display database 66 and screen partition database 62. The  
25 screen-layout-database 50 comprises a number of possible layouts for a monitor screen. In particular it enumerates the locations of the rectangular regions that the screen has been partitioned into. Associated with each partition  
30 are a number of attributes, such as, GRID present

-49-

(is a grid to be drawn onto the partition),  
SELECTABLE (is this partition selectable by the user  
using the mouse), TRANSFORMABLE (can images in this  
partition be transformed or otherwise manipulated),  
5 etc. Additionally, this structure contains the  
locations for graphic or textual objects that will  
be placed by default onto the display screen. For  
instance, the location where the name of the image  
is to be shown, the location where the orientation  
10 of the image (top, left, etc.) is to be shown, etc.  
There may be numerous screen-layout records each  
corresponding to a different layout of the screen.  
Each screen-layout record has a unique ID and a  
textual description.

15 The image database 60 catalogs each image  
within the workstation 15 by a unique identification  
number assigned to the image upon retrieval from the  
image archive/processor 17. The image database 60  
holds a record for each image in the workstation 15,  
20 each record describing an image. In particular each  
record describes:

- 1) the physical-characteristic-of-the-image  
including the number of columns and rows in the  
image, the number of bits in each pixel, the  
25 type of pixel data (integer, floating), the  
number of data values that make up each pixel  
(1 number for REAL pixels, 2 numbers for  
IMAGINARY, etc.);
- 2) the default-display-attributes for the  
30 image, that is how the image is initially to be

-50-

displayed in a partition. It lists attributes such as ZOOM (what magnification is to be applied to the image), ORIENTATION (what is the default gray scale transformation to be applied to the image), etc.

5

3) the image-owner's-demographics including what the image is, what it is about, and how was it acquired. For example, the image is a computed-tomogram of the brain, the image is the sixth in a series of twenty of the patient John Doe, and that John Doe is a 28 year old male with Social Security number 123-23-1234.

10

An index tree may be used to provide conversion between the image identification number and the system wide known image name.

15

The display database 66 comprises a number of records, one for each kind of different hardware display (i.e. monitor) known to the workstation. This database 66 describes the pixel resolution of the display monitor (number of rows and columns), the bit depth of the display monitor (1 bit, 8 bits, etc.), the capabilities of the display monitor (monochrome, color, hardware zoom functions, etc.), and pointers to the low-level support software for this display monitor. The low-level routines are a consistent set of routines for all displays that provide for very simple display functions such as initialize the display monitor, load a rectangular region of display memory with the specified data, etc.

20

25

30

The above three databases the Screen-layout 50, the Image 60 and the Display 66 are primarily

-51-

read-only and are not usually modified by the image display software or the application programmer.

The screen-partition database 62 comprises a number or records, one for each partition on each screen of the monitors 23. By manipulating the contents of this database 62 the user, through the application programmer, causes the desired images to be manipulated and processed. Each record of this database 62 has the following parts:

- 1) an IMAGE-ID-NUMBER which uniquely identifies throughout the system the image in the screen partition corresponding to this record;
- 2) a CURRENT-DISPLAY-ATTRIBUTES which indicate how the image is currently to be displayed on the corresponding screen partition. This includes attributes such as ZOOM, ORIENTATION, GREY-SCALE-TRANSFORMATION, and CENTER (where the center of the image is in relation to the center of the partition). This part is formed of a number of fields each describing a particular display attribute.
- 3) a HIGH-LEVEL-TRANSFORM which describes the next high-level function to be performed upon this image, such as a filter transform, edge enhancement, image smooth, etc. In addition this part holds any of the necessary parameters necessary to further specify the transform, for example, filter coefficients.

30

-52-

4) COMMANDS-TO-EXECUTE that indicates which of the current-display-attributes fields have been modified.

5) OVERLAY-GRAPHICS which describes the graphic primitives currently on the image and that have been drawn upon the image by the user or by system utilities.

6) NON-COMMUTING TRANSFORMS which lists the non-commuting or order dependent transforms applied to the image. That is, for some image transforms the order in which they have been applied to an image is important. That is, for two non-commuting image transforms F and G applied to image 1, we have  $F(G(1))$  does not equal  $G(F(1))$ . This part of data structure 62 represents a time ordered list of the non-commuting transforms applied to the image, their associated parameters, and the current display attributes.

Thus, the screen partition database 62 provides a compressed and complete description of all images on all monitors in the workstation 15.

Referring now to Figure 11, the major program modules or image display software routines are described next. The module DO-TRAVERSE-S-P-DATABASE 54 is responsible for traversing the screen-partition database 62 to search for any COMMANDS-TO-EXECUTE or HIGH-LEVEL-TRANSFORMS needing to be performed for a user request. For each image processing command found, the module 54 decides if

-53-

the command can be done locally (i.e., within the work-station 15) or must be performed by the shared image archive/processor 17. For each command to be performed locally, the module 54 passes the request to the module DO-IMAGE-FUNCTION 58. For each command that is to be performed by the shared image archive/processor 17, module DO-TRAVERSE-SCREEN-PARTICIAN-DATABASE 54 makes two requests through modules HEAD-END-COMMUNICATION 74 and IMAGE-DATA-CACHE 56. A first request involves application of the particular command to the specified image. The second request asks that the resultant image be transmitted to the work-station 15 and module PENDING-MONITOR 52 be informed that an image with a specified ID is expected, and furthermore, upon its arrival is to be placed at a specified screen and partition (via the command IMAGE-LOAD). Also the module DO-TRAVERSE-S-P-DATABASE 54 clears the request for the command from the COMMANDS-TO-EXECUTE or the HIGH-LEVEL-TRANSFORM field of the Screen Partician Database 62.

The module DO-IMAGE-FUNCTION 58 collects all the relevant data necessary to perform the function on the image at the designated screen and partition, as requested by the module DO-TRAVERSE-S-P-DATABASE 54. This includes gathering the location of the partition, the physical attributes of the image, and pointers to the image data in the local cache/working buffer 35. If all the data is in the workstation 15, specifically the image data is in



-54-

the local cache/working buffer 35, this module 58 passes the request command along with the associated supporting image data to the module IMAGE-PRIMITIVE 64 for processing. Otherwise, its work is  
5 completed.

The module IMAGE PRIMITIVE 64 chooses the appropriate display-specific-low-level-driver, converts the relatively high-level requests received from DO-IMAGE-FUNCTION 58 to the low-level data  
10 abstractions needed by the functions of the display-specific low-level drivers, decides which low-level functions are needed to perform the requested function, and then invokes these functions in the display-specific low-level driver. Each  
15 display specific low-level-driver supports a collection of low level functions of various types that perform for each type of hardware display a uniform set of functions such as, initialize the display, clear the display, and load a block of  
20 pixel data into a rectangular section of display memory. There is one collection of low level functions for each type of hardware display connected to the workstation 15.

The module IMAGE DATA CACHE 56 manages a block  
25 of local high-speed storage that is used as a local cache for images. When pointers to image data are requested by module DO-IMAGE-FUNCTION 58 the IMAGE DATA CACHE module 56 either 1) if the image data is present, returns pointers to module DO-IMAGE-

30

-55-

FUNCTION 58, or 2) if the image data is not present, passes a request to module HEAD-END-COMMUNICATIONS 74. In turn module 74 requests the image archive/processor 17 to (i) send the needed image data, and (ii) inform module PENDING-MONITOR 52 that an image with a specified ID is expected and that upon its arrival is to be placed at a specified screen and partition with the original desired function.

10       The module HEAD-END-COMMUNICATIONS 74 manages requests for image archive/processor 17 support. For incoming images, HEAD-END-COMMUNICATIONS module 74 requests module IMAGE-DATA-CACHE 56 to store images and then informs module PENDING-MONITOR 52 of  
15       the image arrival. Additionally, this module 74 passes requests to module PENDING-MONITOR 52 from module IMAGE-DATA-CACHE 56.

          The module PENDING-MONITOR 52 logs multiple messages specifying that soon images will be  
20       received from the image archive/processor 17 that are intended to be placed upon specific screens and partitions. Upon receiving from module HEAD-END-COMMUNICATIONS 74 a message that an image with a particular ID has arrived, module PENDING-MONITOR 52  
25       searches within an internal data structure for the message that indicates what screen and partition the image is to be displayed upon. Upon finding the associated message, module PENDING-MONITOR 52 inserts either the original command or the command  
30       to load the specified image upon the indicated

-56-

screen and partition into the screen-partition database. The PENDING-MONITOR 52 then invokes module DO-TRAVERSE-S-P- DATABASE 54 to process the function/command as though the function was being  
5 invoked for the first time directly by a user request.

Interaction between the foregoing program modules for image processing most commands typically starts with the application programmer changing the  
10 image ID number, current display attributes, and/or high level transforms in the screen partition database 62 to effect the desired function. This in effect records any changes needed on the monitors 23 or in the user desired images. The commands-to-  
15 execute field is modified either explicitly by the application programmer or indirectly by the process of modifying the other three fields. After indicating all changes to the display screens 23 through changes to the screen-partition database 62  
20 the applications programmer has only to invoke the module DO-TRAVERSE-SCREEN-PARTITION-DATABASE 54 to cause the display screens 23 to be updated to reflect the database 62. Thereafter one of three main control flow paths is taken through the main  
25 support program 80 of Figure 11 depending on if the image data is present in the cache/working buffer 35 and if the request image processing command can be performed locally or must instead be performed by the image archive/processor 17.

30

-57-

In the first case where the function or command can be done locally and the image data is in the cache/working buffer 35, DO-TRAVERSE-S-P-DATABASE 54 passes the request to module Do-Image-Function 58 and clears the command from the screen participant database 62. DO-IMAGE-FUNCTION 58 requests and receives from the IMAGE-DATA-CACHE 56 pointers to the image data in the cache/working buffer 35. After collecting the pointer to the image data and all necessary information (such as coordinates of the destination screen and partition, description of the image x-y pixels, bit depth, etc.), the DO-IMAGE-FUNCTION module 58 passes the collected information and an indication of the destination partition to image primitive module 64. In turn IMAGE-PRIMITIVES-64 determines which DISPLAY-SPECIFIC-LOW-LEVEL-FUNCTIONS are needed to perform the requested functions and which driver to invoke. IMAGE-PRIMITIVE 64 thereafter invokes the necessary routines in the appropriate driver. The chosen low-level functions move, and if necessary transform, the pixel data to the display buffers to modify the monitors as desired.

In a second case, DO-TRAVERSE-S-P-DATABASE 54 determines that the command can be serviced locally and passes the request to DO-IMAGE-FUNCTION 58 which requests the image data from the image cache/working buffer 35 via IMAGE-DATA-CACHE 56. Finding the image data is not in the cache 35, DO-IMAGE-FUNCTION comes to completion. IMAGE-DATA-CACHE 56 passes a

-58-

message to HEAD-END-COMMUNICATION MODULE 74 to request the image archive/processor 17 to send the desired image to the workstation 15. Module 74 transfers the request to image archive/processor 17 and informs the pending monitor 52 that image data is expected from the image archive/processor 17 and is intended to be loaded into a specified screen and partition with the original commands.

After some unspecified time module HEAD-END-COMMUNICATION 74 receives the image and thereafter 1) requests module IMAGE-DATA-CACHE 56 to place the received image into the image cache/working buffer 35 and 2) informs the module PENDING-MONITOR 52 that an image with a specified ID has arrived into the cache 35. Module PENDING-MONITOR 52 looks up the image ID in its internal data structure to determine which screen and partition the image is to be placed upon. Module PENDING-MONITOR 52 then modifies the appropriate screen-partition database record and indicates which command is to be performed. The PENDING-MONITOR 52 then invokes the module DO-TRAVERSE-S-P-DATABASE 54. The remaining processing is as in the first case since now there is a command that can be serviced locally and the image data is now in the local image cache 35.

In the case where the function cannot be processed locally, as determined by the module DO-TRAVERSE-S-P-DATABASE 54, the following occurs.

-59-

By definition the image data cannot be in the local cache 35 at this time and the original image is at the image archive/processor 17. Module DO-TRAVERSE-S-P-DATABASE passes the request on to module HEAD-END-COMMUNICATION 74 requesting that the image archive-processor 17 perform the desired function on the specified image and then send the transformed image to the workstation 15.

Module HEAD-END-COMMUNICATION 74 sends the request to the image archive/processor 17 and informs module PENDING-MONITOR 52 that soon an image will be received and that image is intended to be loaded upon a specified screen and partition.

After some unspecified time module HEAD-END-COMMUNICATION 74 receives the transformed image and 1) requests module IMAGE-DATA-CACHE 56 to place the transformed image into the local image cache 35 and 2) informs module PENDING-MONITOR 52 than an image with a specified ID has arrived into the local cache 35. Module PENDING-MONITOR 52 looks up the image ID in its internal data structure to determine which screen and partition the image is to be placed upon. The module PENDING-MONITOR 52 next modifies the appropriate screen-partition database record and indicates that the command is an IMAGE-LOAD. The PENDING-MONITOR 52 then invokes the module DO-TRAVERSE-S-P-DATABASE 54. The remaining processing is as in the first case since now there is a command that can be serviced locally (IMAGE-LOAD) and the image data is now in the local image cache 35.

-60-

While the invention has been particularly shown and described with reference to a preferred embodiment thereof, it will be understood by those skilled in the art that various changes in form and details may be made therein without departing from the spirit and scope of the invention as defined by the appended claims. For example, in a particular application of the image/graphics system of Figures 1a-1b, radiological image data such as x-rays and magnetic resonance images are stored in an image archive 17 and are transferred to workstations 15 on the one-way optical bus 31 upon user request. Graphics data, such as text and geometric markings (lines, boxes, etc.) provided and overlaid on a displayed image by a doctor during diagnosis, are passed to the workstations 15 via the general purpose bus 21.

15.

-61-

CLAIMS

1. An image display system for displaying images stored in a remote source, comprising:

5 a plurality of display means viewable by a user;

10 a plurality of display controllers, one display controller connectable to one display means at a time to drive the display means, each display controller providing to a display means image data indicative of an image to be displayed, at least one display controller being unconnected to any display means while all display means are being driven by the other display controllers; and

15 means for changing connections between display controllers and display means such that an unconnected display controller becomes connected to a display means and a display controller currently connected to the display means becomes disconnected, the display means thereby being driven in a paint-free fashion from image data from the currently connected display controller and then from image data from the unconnected display controller.

- 20 2. An image display system as claimed in Claim 1 wherein there is one more display controller than display means.
- 25



-62-

3. An image display system as claimed in Claim 1 wherein the means for changing connections include a cross-bar switching circuit having a plurality of outputs connectable to the display means in a one to one correspondence and a plurality of inputs connectable to the display controllers in a one to one correspondence, each input associable with a different output; the circuit enabling throughput of image data from all but at least one input to their respectively associated outputs.
4. An image display system as claimed in Claim 3 further comprising:
- a plurality of video mixers one for each display means, each video mixer connected between a different display means and a respective output of the cross-bar switching circuit, such that image data indicative of the images from all but at least one of the display controllers are transferred to the video mixers; and
- means for providing to the video mixers graphic signals indicative of user generated data associated with the images provided by the display controllers; each video mixer combining respective image data and graphic signals of an image and driving the connected display means therewith to provide a display of the image overlayed with the user generated data.

-63-

5. An image display system as claimed in Claim 4 wherein:

the means for providing graphic signals to the video mixers includes a general purpose bus for receiving the user generated data; and

the display controllers receive image data indicative of images to be displayed from a bus which is dedicated to carrying image data and which is independent of the general purpose bus.

6. An image display system as claimed in Claim 1 further comprising an image data bus for passing image data to the unconnected display controller such that when the display controller becomes connected it drives a display means in a manner that displays an image on a user specified partition of the display means.

7. An image display system as claimed in Claim 1 further comprising an image data bus for passing image data between the display controllers, the bus enabling copying of image data of an image from a display controller which is connected to a display means to a display controller which is unconnected to any display means while the connected display controller continues driving the display means.

-64-

8. An image display system as claimed in Claim 7 further comprising:

5 a cache memory for locally holding image data indicative of images to be displayed, the cache memory receiving image data from the remote source;

10 the image data bus enabling transferring of image data from the cache memory to a display controller which is unconnected to any display means.

9. An image display system as claimed in Claim 8 wherein the cache memory includes means for providing a grey scale transformation for each image to be transferred from the cache memory, the type of grey scale transformation being changeable for each transfer of image data from the cache memory on the image data bus . .

10. An image display system as claimed in Claim 8 wherein the cache memory includes means for modifying size of an image to be displayed, the size of an image being changeable for each transfer of image data from the cache memory on the image data bus.

11. An image display system as claimed in Claim 1 wherein each display controller includes rectangle loader means for displaying an image at least on a first user specified portion of a first display means.

-65-

12. An image display system as claimed in Claim 1 further comprising two communication channels, one channel for primarily transferring image data to the system from the remote source and the other channel connected to computer means for carrying control signals necessary to enable the transferring of image data on the one channel.
13. An image display system as claimed in Claim 12 further comprising receiving means for decompressing image data being transferred from the remote source to the system, the receiving means changeably using one decompression scheme on a part of the image data and another decompression scheme on a remaining part of the image data to decompress the image data.
14. An image display system as claimed in Claim 1 wherein the display means include display monitors for displaying medical images.
15. An image display system for operation in conjunction with computer means and a remote source of images which are selectively displayed, the system comprising:  
a plurality of display means for displaying user selected images, each display means having a screen from which a user views the images;

-66-

processing means for receiving user  
selected images from the remote source and for  
driving a corresponding user selected display  
means to display the selected image;

5        an image data channel dedicated to  
transferring image data indicative of the  
selected images from the remote source to the  
processing means; and

10        a host channel coupled to the computer  
means for providing command and control signals  
to the processing means to enable the  
transferring of image data from the remote  
source on the image data channel to the  
processing means.

15    16. An image display system as claimed in Claim 15  
wherein the image data channel is  
unidirectional and has a relatively high  
bandwidth.

20    17. An image display system as claimed in Claim 15  
wherein the host channel provides and receives  
command and control signals to and from the  
processing means in a manner that enables the  
image data channel to transfer image data  
substantially free of such command and control  
25    signals.

-67-

18. An image display system as claimed in Claim 15 wherein the processing means includes:

5 a general purpose bus for transferring command and control signals throughout the processing means;

10 a plurality of display controllers, one each for receiving image data indicative of a selected image and driving a corresponding display means therewith; and at least one controller remaining free from driving a display means;

15 an image bus for transferring image data received from the image data channel to the display controllers and for transferring image data between the display controllers;

an image processing bus for transferring image data being processed within the processing means; and

20 means for switching a remaining display controller with a display controller driving a display means in a manner which provides an abrupt transition from one selected image to another as displayed on the corresponding display means, the display controller initially driving the display means becoming a display controller remaining free from driving a display means, the image bus in a relatively high bandwidth manner providing image data to the display controller remaining free from driving a display means such that the means for switching is able to provide an abrupt transition on any of the display means.

25

30

-68-

19. An image display system as claimed in Claim 15 wherein the means for processing further comprises means for receiving image data from the image data channel and for decompressing the image data with different schemes according to high and low order parts.
20. An image display system, comprising:
- a. a source of images having computer means;
  - b. a plurality of independent workstations having access to and sharing the source of images, each workstation having:
    - i) a multiplicity of display means for displaying user selected images from the source of images; and
    - ii) processing means for receiving user selected images from the source and for driving a user selected display means within the workstation to display the selected images; and
  - c. channel means for providing image data to the workstations from the source of images, and for coupling the source and workstations together to enable processing of image data by the processing means and computer means of the source of images in a manner free of transmission of image data from the workstations to the source, such that receiving and processing of images in each workstation appears to be independent and interference free of receiving and processing of images in the other workstations.

-69-

21. An image display system as claimed in Claim 20 wherein the computer means of the source includes an image cache for holding images anticipated to be transferred from the source to the workstations.
22. An image display system as claimed in Claim 20 wherein the computer means of the source employs assignment of priorities to processing tasks such that frequently desired tasks are given higher priority than computationally intensive tasks.
23. An image display system as claimed in Claim 20 wherein the processing means includes:
- a first routine for passing to the computer means of the source each user image processing request which includes a function unable to be performed by the processing means; and
  - a second routine for subsequently receiving from the computer means image data generated from the user image processing request as performed by the computer means and for loading the image data into a user desired display means.



-70-

24. An image display system as claimed in Claim 20 wherein the processing means includes:

5 a first routine for passing to the computer means of the source each user command to process an image which is remote from the workstation, and

10 a second routine for subsequently receiving from the computer means the image and for locally processing the user command with the received image.

25. An image display system for displaying images selected from a remote source, the system comprising:

15 at least one display means; processing means for driving the display means with user desired images selected from the remote source and stored in local memory of the processing means;

20 an image data base for indicating original attributes including dimensions of each selected image retrieved from the remote source and subsequently stored in the local memory;

25 means for partitioning each display means in a user desired manner; and

-71-

5 a display data base having a record for  
each partition of each display means, each  
record providing, for a corresponding partition  
of a display means, an indication of a user  
selected image to be displayed therein with  
respective user desired display attributes  
thereof, the processing means using the display  
data base to drive the display means such that  
10 the user desired images selected from the  
remote source are displayed in user desired  
partitions of the display means with the user  
desired attributes.

26. An image display system as claimed in Claim 25  
wherein the processing means comprises:

15 an initial routine for changing the  
display database upon a user request for a  
desired image process to be performed on a  
desired image forming a processed image, the  
initial routine changing the display database  
20 to indicate the desired image and processor  
function associated with the desired process;  
and

a traverse routine for subsequently  
reading the display database in response to the  
25 changing of the display database, the traverse  
routine causing the processor function to be  
performed and thereafter driving the display  
means such that the requested processed image  
is displayed on a user desired display means.  
30

-72-

27. An image display system as claimed in Claim 25 wherein:

the remote source includes computer means;  
and

5 the processing means includes:

a first routine for passing to the  
computer means of the source each user image  
processing request which includes a function  
unable to be performed by the processing  
10 means; and

a second routine for subsequently  
receiving from the computer means image data  
generated from the user image processing  
request as performed by the computer means and  
15 for loading the image data into a user desired  
display means.

28. An image display system as claimed in Claim 25 wherein:

the remote source includes computer means;  
20 and

the processing means includes:

a first routine for passing to the  
computer means of the source each user command  
to process an image which is remote from the  
workstation, and  
25

a second routine for subsequently  
receiving from the computer means the image and  
for locally processing the user command with  
the received image.

30

29. An image display system as claimed in Claim 25 wherein, in response to a user request to copy an image from an initial partition of one display means to a different partition of another display means, the processing means copies the image indicated in the record corresponding to the initial partition into the record corresponding to the different partition and thereafter drives the display means according to the records of the display database.
30. An image display system as claimed in Claim 25 wherein, in response to a user request to copy a whole screen view of one display means to a second display means, the processing means:
- (i) through the means for partitioning, changes partitioning of the second display means to match partitioning of the one display means;
  - (ii) copies the records corresponding to the partitions of the one display means into the respective records corresponding to the matching partitions of the second display means; and
  - (iii) drives the display means according to the records of the display database.

-74-

31. An image display system as claimed in Claim 25 further comprising means for saving a display status of the system, the means including:  
an indication of the partitioning of each display means; and  
a copy of the display database.
32. An image display system as claimed in Claim 25 wherein each record of the display database includes an ordered list of non-commuting transforms applied to the image to be displayed in the corresponding partition, the list of non-commuting transforms together with the display attributes providing an indication of formation of the image.
33. An image display system as claimed in Claim 25 wherein each record of the display database includes an indication of user selected graphics overlayed on the image to be displayed in the corresponding partition.
34. An image display system as claimed in Claim 25 wherein the image database includes for each selected image:  
a first field for indicating physical characteristics of the image;  
a second field for indicating default display attributes; and  
a third field for indicating image owner information.

-75-

35. An image display system as claimed in Claim 25 wherein the means for partitioning include a screen layout database for indicating a multiplicity of possible screen layouts including partitions of a display means, the screen layout database having a record for each possible screen layout, each record including:
- an identification field for textually describing the respective screen layout;
  - a count field for indicating the number of partitions involved in the screen layout;
  - a location and attribute field indicating location and attributes of each partition in the screen layout; and
  - a graphics field indicating default location of non-image data to be displayed in the screen layout.
36. An image display system as claimed in Claim 25 wherein each record of the display database includes:
- an image identifier which uniquely identifies the selected image to be displayed in the partition of the display means corresponding to the record;
  - an attributes element for indicating how the image is to be currently displayed, the attributes element having an attribute field for each possible attribute;
  - a transform field for indicating a high level function to be currently performed on the image;

-76-

a changed-attribute field indicating attribute fields which have been modified;  
a graphics field for indicating graphics which are currently overlayed on the image; and  
5 a list of non-commuting transforms which have been applied to the image.

37. An image display system as claimed in Claim 25 further comprising a display means database for describing the display means, the display means  
10 database having a record for each display means, each record including:

a resolution component for indicating pixel resolution, spatial resolution and bit depth of the respective display means;  
15 a function field for indicating possible physical functions of the display means; and  
a pointer to low level processor support means.

38. An image display system for displaying images  
20 selected from a remote source of compressed images, the image display system comprising:  
an image channel for receiving from the source image data indicative of the images, the image data being in the form of a series of  
25 data words, each data word having a sequence of high order parts and a corresponding sequence of low order parts; and

decompression means for decompressing image data received through the image channel,  
30 the decompression means decompressing high order parts and corresponding low order parts of a series of different data words with different methods.

-77-

39. An image display system as claimed in Claim 38 wherein the decompression means decompresses high order parts of the series of different data words by Run Length Coding when the high order parts within the series change less often than the corresponding low order parts.
40. An image display system as claimed in Claim 38 wherein the the decompression means decompresses the low order parts of the series of different data words by Delta Coding.
41. An image display system as claimed in Claim 38 wherein the decompression means decompresses the high order parts of the series of different data words according to number of similar high order parts in immediate succession in the series.
42. An image display system as claimed in Claim 41 wherein the decompression means decompresses the high order parts of the series by:
- Run Length Coding - Long method, where the number of similar high order parts in succession is 13 or more;
  - Run Length Coding - Short method, where the number of similar high order parts in succession is at least 11;
  - Delta - 2 Coding method, where the number of similar high order parts in succession is at least 8; and
  - Delta - 4 Coding method, where the number of similar high order parts in succession is at least 6.



-78-

43. An image display system as claimed in Claim 38 wherein the image data received from the source comprises:

5 a series of compression units, each compression unit being preceeded by a respective compression unit header which identifies the compression unit; and

10 an image header identifying the number of compression unit headers within the data storage element.

44. An image display system as claimed in Claim 43 further comprising:

15 a buffer memory for locally holding image data received from the source and for subsequently providing said data to the decompression means;

20 processing means for transferring on the image channel whole compression units from the source to the buffer memory in a compression unit by compression unit manner.

45. An image display system as claimed in Claim 43 further comprising a segmented buffer memory for locally holding decompressed image data resulting from decompression of compression units of image data received from the source, the buffer memory having a plurality of segments, resulting decompressed image data from one compression unit being held in one segment of the buffer memory for each compression unit.
- 25
- 30

-79-

46. An image display system for displaying images stored in a remote source, comprising:

a processor;

a plurality of display means

5 simultaneously driven by the processor to display user selected images from the remote source, one display means providing both an indication of images currently being displayed on the other display means and an indication of  
10 user selected images to be displayed on the other display means; and

image control means for processing through the processor user selections of images to be displayed and user selections of display means  
15 on which each image is to be displayed.

47. An image display system as claimed in Claim 46 wherein the images include medical images.

48. In an image display system having a plurality of display means, a method for displaying  
20 images stored in a remote source, comprising the steps of:

providing image data to a plurality of display controllers, one display controller connectable to one display means at a time to  
25 drive the display means with the image data, each display controller providing to a display means image data indicative of an image to be displayed, at least one display controller being unconnected to any display means while  
30 all display means are being driven by the other display controllers; and

-80-

changing connections between display controllers and display means such that an unconnected display controller becomes connected to a display means and a display controller currently connected to the display means becomes disconnected, such that the display means is driven in a paint free fashion from image data from the currently connected display controller and then from image data from the unconnected display controller.

49. A method for displaying as claimed in Claim 48 wherein the step of changing connections includes providing an analogue crossbar switching circuit having a plurality of outputs connectable to the display means in a one-to-one correspondence and a plurality of inputs connectable to the display controllers in a one-to-one correspondence, each input associable with a different output in a manner which enables throughput of image data from all but at least one input to respectively associated outputs.

50. A method for displaying as claimed in Claim 48 further comprising the steps of:

providing graphic signals indicative of user generated data associated with the images provided by the display controllers; and combining respective image data and graphic signals of an image and driving the connected display means therewith to provide a display of the image overlaid with the user generated data.

-81-

51. A method for displaying as claimed in Claim 48 wherein the step of providing image data to a plurality of display controllers includes transferring image data between display  
5 controllers and between a local memory and the display controllers while the connected display controllers continue driving the display means.
52. A method for displaying as claimed in Claim 48 further comprising the step of:  
10 in the unconnected display controller, preparing image data to be subsequently displayed.
53. A method for displaying as claimed in Claim 52 wherein the step of preparing image data in the  
15 unconnected display controller includes transferring image data from a connected display controller to the unconnected display controller.
54. A method for displaying as claimed in Claim 52 wherein the step of preparing image data in the  
20 unconnected display controller includes transferring image data from a local memory area to the unconnected display controller.

-82-

55. In an image display system, having a plurality of display means for displaying user selected images, a method of displaying images selected from a remote source, comprising the steps of:
- 5        providing processor means for receiving user selected images from the remote source and for driving a corresponding user selected display means to display the selected image;
- 10        transferring image data indicative of the selected images from the remote source to the processing means on an image only carrying channel; and
- 15        providing on a host channel command and control signals between a host processor and the processing means to enable the transferring of image data from the remote source on the image only carrying channel to the processor means.
- 20        56. A method for displaying images as claimed in Claim 55 wherein the step of providing processor means includes providing means for receiving image data from the image only carrying channel and for decompressing the image data with different schemes according to
- 25        high and low order bits..

-83-

57. In a network of independent workstations, each workstation having a plurality of display means for displaying user selected images from a remote source shared by the workstations, a method of displaying images comprising the steps of:

providing a host computer supporting each workstation;

providing in each workstation processing means for receiving user selected images from the source and for driving user selected display means within the workstation to display the selected images;

transferring image data indicative of the selected images from the remote source to the processing means on a unidirectional image-only carrying channel;

for each workstation, transferring on a host channel between the host computer and the processing means of a workstation command and control signals to enable the transferring of image data from the remote source on the image-only carrying channel to the processing means, and processing the image data locally, with respect to the workstation, using the processing means, such that receiving and processing of images in each workstation appears to be independent and interference free of receiving and processing of images on the other workstations.

-84-

58. In an image display system for displaying  
images selected from a remote source of  
compressed images, the image display system  
having a channel for receiving from the source  
5 image data indicative of the images in the form  
of a series of data words each data word having  
a sequence of high order parts and a  
corresponding sequence of low order parts, the  
method of decompressing image data comprising:  
10 decompressing high order parts and  
corresponding low order parts of a series of  
different data words with different methods.
59. In an image display system for displaying  
images stored in a remote source, the method of  
15 displaying images comprising:  
providing a processor;  
providing a plurality of display means  
simultaneously driven by the processor to  
display user selected images from the remote  
20 source, one display means providing both an  
indication of images currently being displayed  
on the other display means and an indication of  
user selected images to be displayed on the  
other display means; and  
25 processing through the processor, while  
images are being displayed on the display  
means, user selected images to be later  
displayed in user selected partitions of the  
display means.

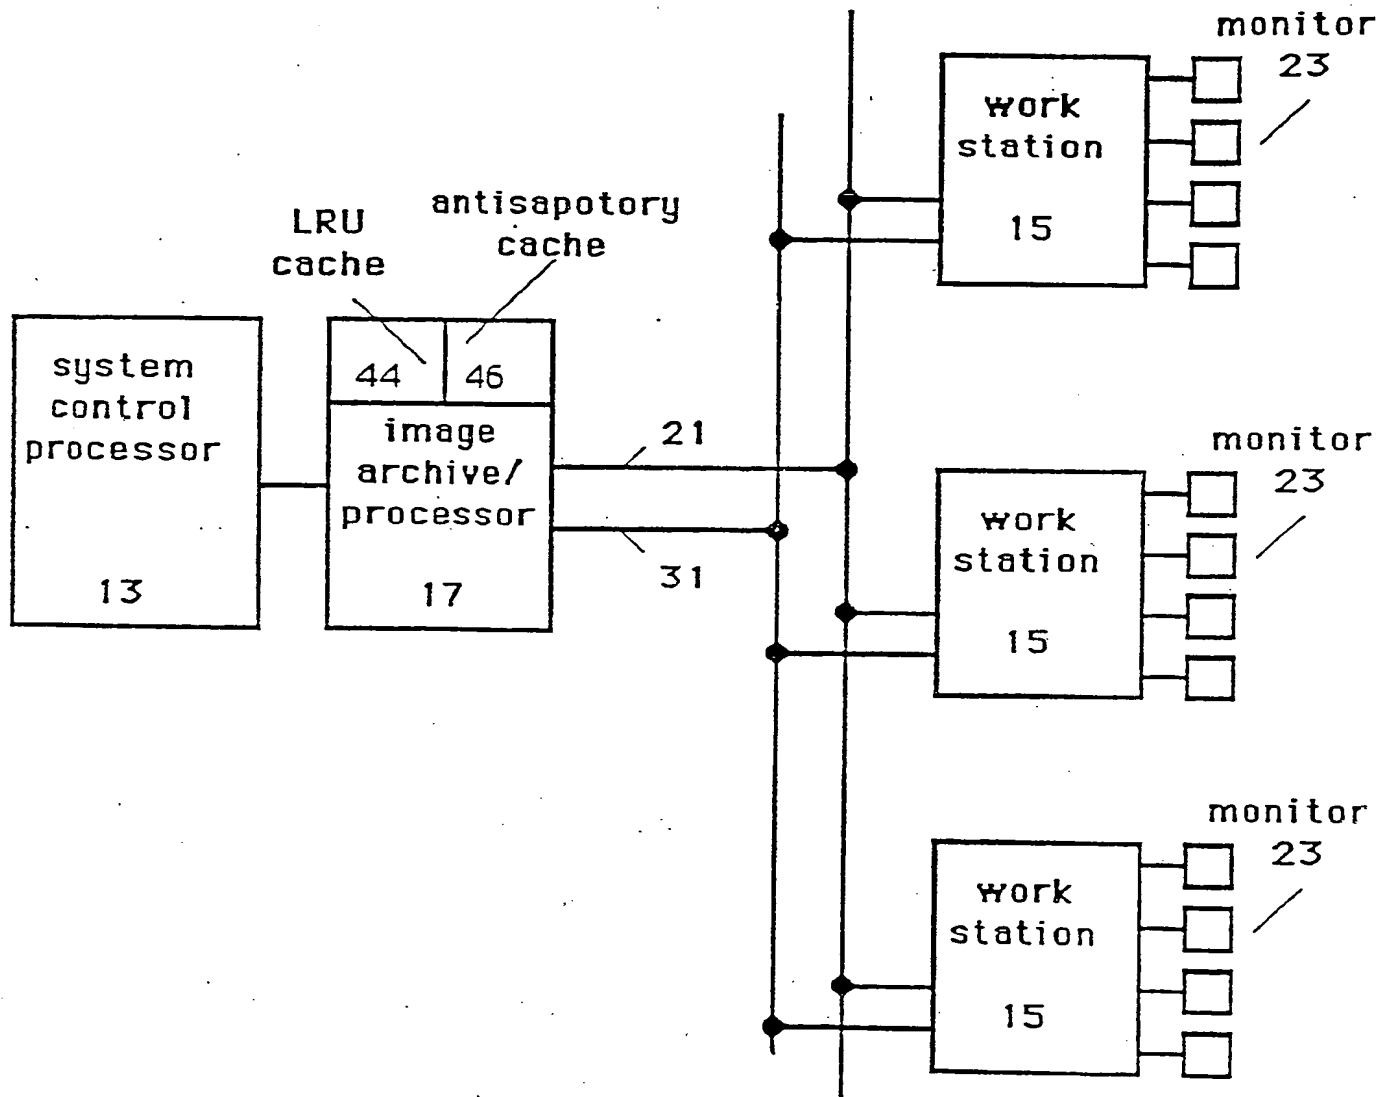


Fig. 1a






Fig. 1b

3/11




Fig. 1c





Fig.


411

Figure 2a  
Analog Cross-Bar Switch.



45 →

5/11

Figure 26, Analog Transfer Gate

6/11

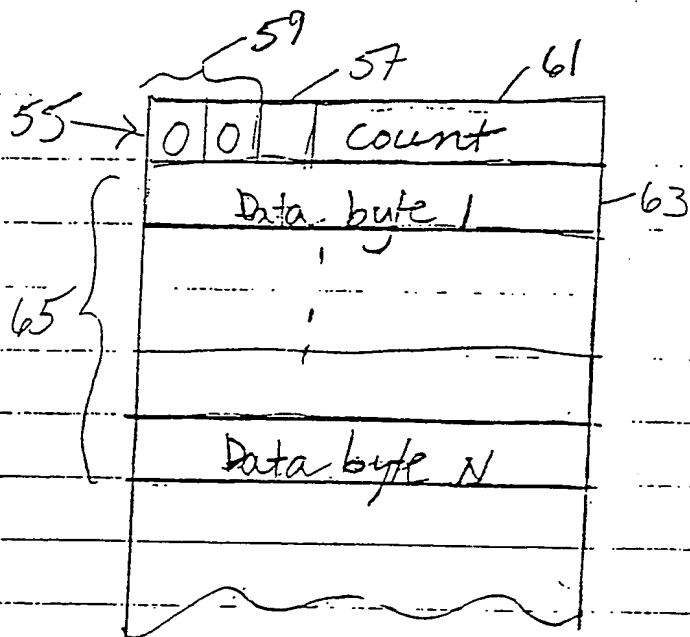


Fig. 3

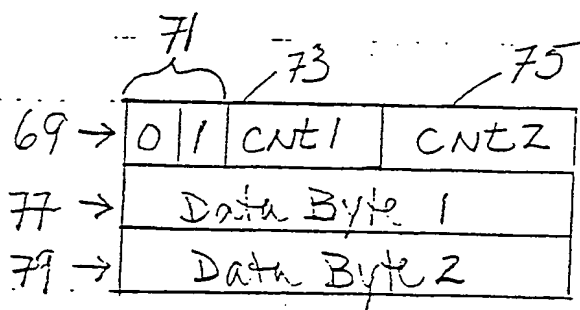


Fig. 4




Fig. 5

7/11

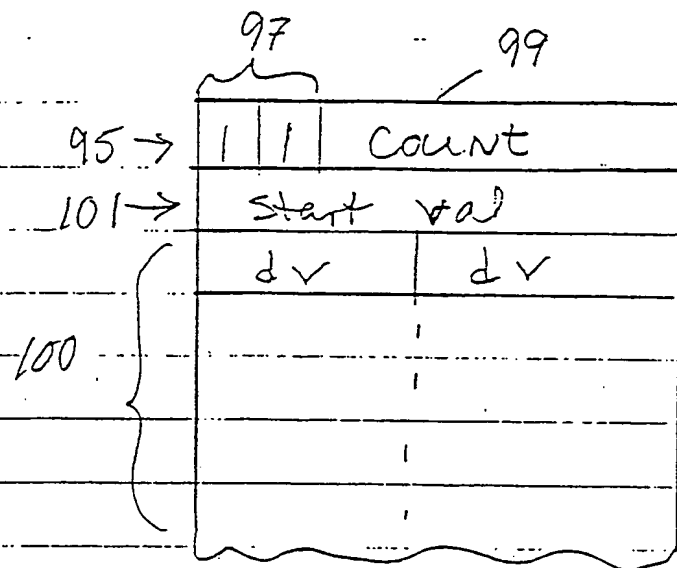


Fig. 6

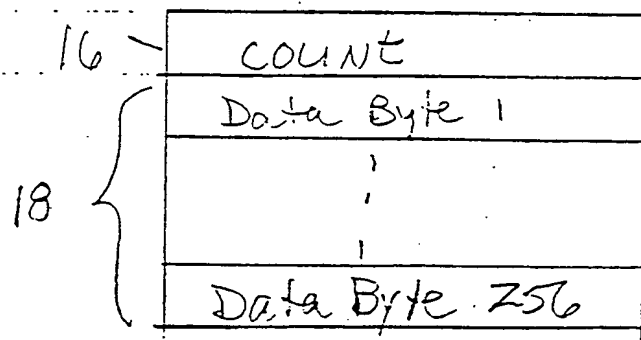


Fig. 7




Fig. 8

8/11

Fig 9.

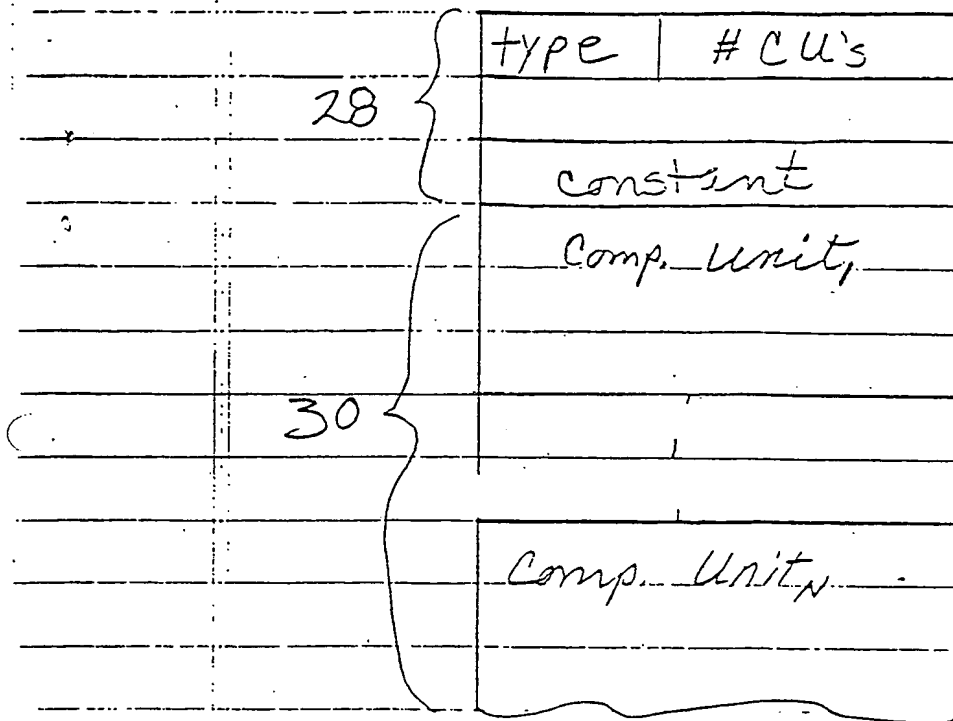

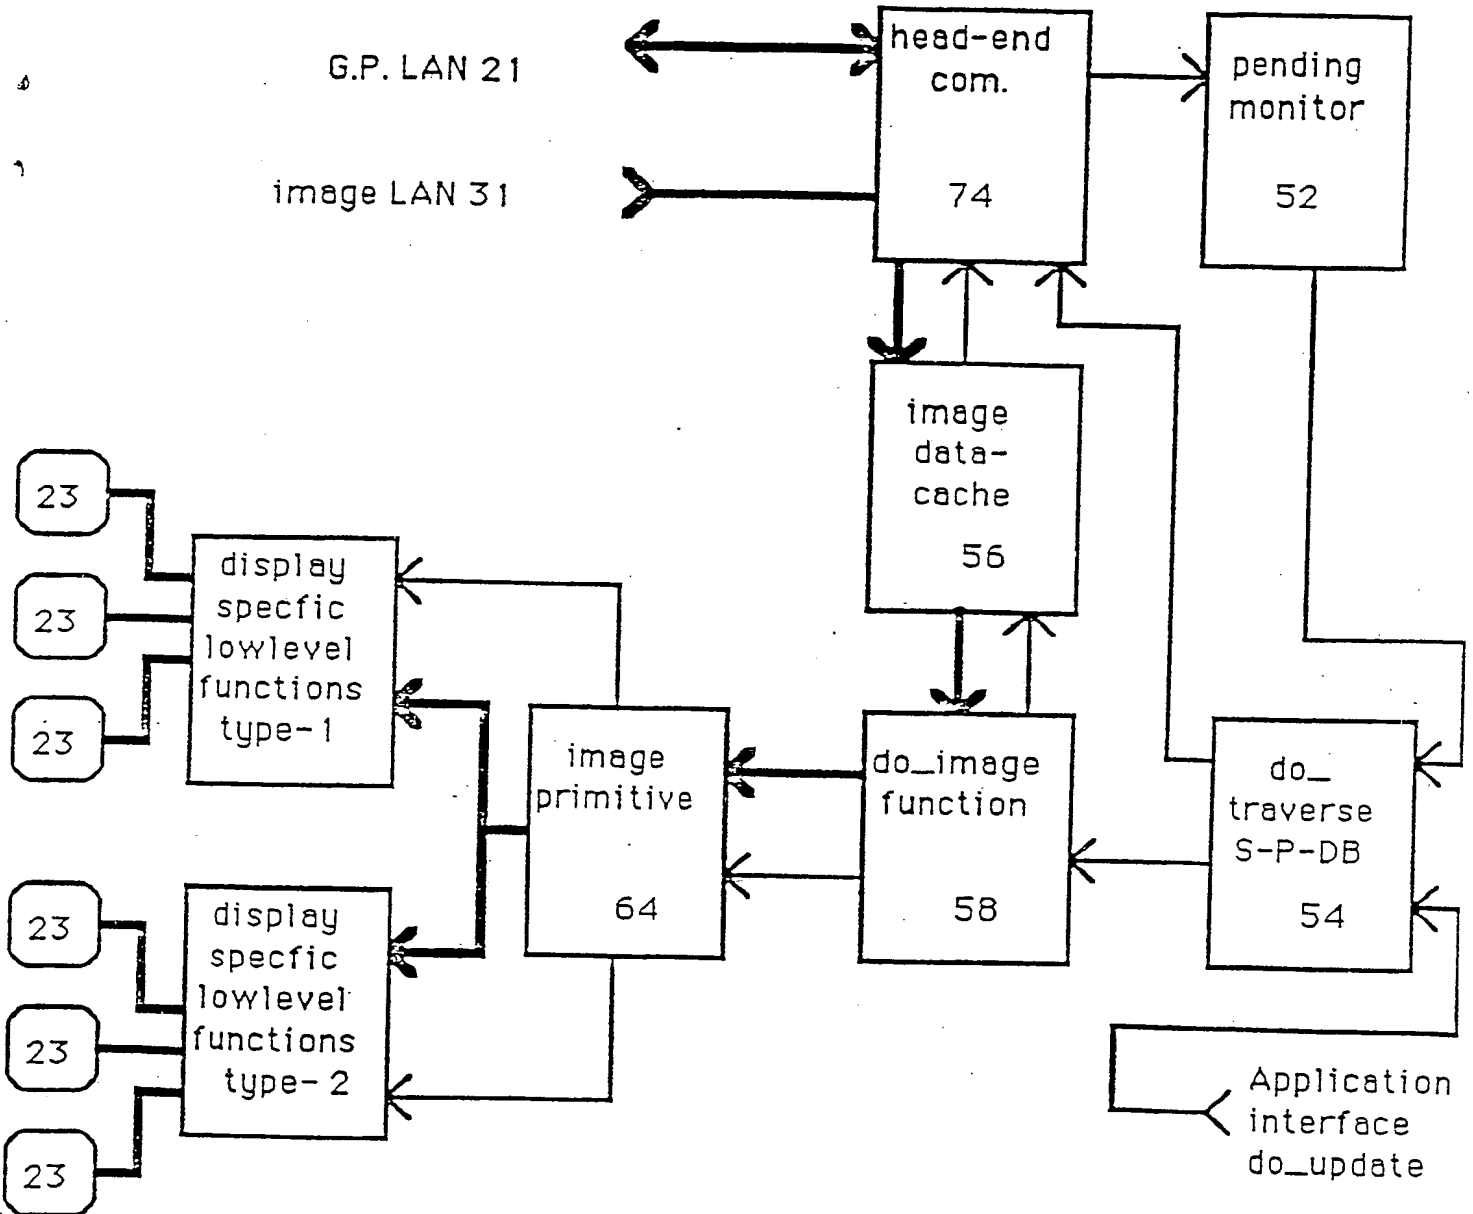



Fig 10



9/11



↑

80

Fig. 11



10/11

Screen layout database
Accessed by Layout ID
text description
number of partition this screen
location of all partitions
attributes of all partitions
graphic object locations

Fig. 12a

↑ 50

Image database
Accessed by Image ID
physical characteristics of image
default display attributes
image owner's demographics

Fig. 12b

↑ 60

///

Display database
Accessed by Display ID
pixel resolution of display
spatial resolution of display
bit depth of display
capabilities of display
pointer to display support software

Fig. 12c

↑ 66

Screen-partition database
Accessed by Screen-partition
pending commands to execute
high level transformations
partition attributes
image ID number
current image display attributes
non-commuting transforms
overlay graphics

Fig. 12d

↑ 62